

E0170

FORCE BITS FOR MAC CONFIGURATION

RELATED APPLICATIONS

This application claims priority from Provisional Application Serial No. 60/173,486 filed on December 29, 1999 entitled: "FORCE BITS FOR MAC CONFIGURATION", the entire disclosure of which is hereby incorporated by reference herein.

5 FIELD OF THE INVENTION

The present invention relates to a local area network interface for a multiple port device, more particularly to an arrangement for configuring physical layer devices of the multiple port devices in order to communicate network data between each physical layer device and corresponding link partners.

10 BACKGROUND ART

Local area networks use a network cable or other network media to link nodes (e.g., workstations, routers and switches) to the network. Each local area network architecture uses a media access control (MAC) enabling network interface device (NID) at each network node to share access to the media.

15 Traditional Ethernet networks (10BASE-T) operate at 10Mb/s Ethernet protocol, as described by IEEE Standard 802.3; the majority of Ethernet interfaces currently operate at this data rate. However, a newer Ethernet standard, under IEEE standard 802.3u, accomplishes the faster operation of 100 BASE-T systems, at a 100 Mb/s data rate (i.e., a 125Mb/s encoded bit rate) using unshielded twisted pair (UTP) physical media. The 100 BASE-T standard defines operation over two pairs of
20 category 5 UTP (100 BASE-TX) and over four pairs of category 3 UTP. The 100 BASE-FX network medium, covered by the 100 BASE-T standard, allows operation over dual fiber optic cabling.

Physical (PHY) layer devices are used to enable the MAC to send and receive digital packet data across an interface that is generic to the different media standards. In particular, PHY devices are configured for translating digital packet data received from a MAC across a standardized interface,
25 e.g., a Media Independent Interface (MII), into an analog signal for transmission on the network medium, and reception of analog signals transmitted from a remote node via the network medium. Hence, the MII connects the MAC to a physical layer (PHY) transceiver configured for a particular network medium, e.g., 10 BASE-T, 100 BASE-FX, or 100 BASE-TX.

Newer PHY devices are designed to operate at different network speeds, depending on the
30 type of network medium and the capabilities of a link partner at the other end of the network medium.

In such a case, each PHY device needs to be configured to a setting that is compatible with the configuration settings of the link partner's PHY device.

Auto-negotiation is performed between PHY devices of link partners. For example, auto-negotiation is performed as a link startup procedure between a network switch port (or repeater port) and a workstation linked to the switch port via the network medium each time a link to the switch port is connected, powered on or reset. During auto-negotiation, the two PHY devices of the link partners exchange information about their capabilities, and then the PHYs configure themselves to the best operating mode that is common to them.

Hence, the PHY transceivers of the link partners will typically perform auto-negotiation amongst each other to determine the best common operating mode for transmission and reception of data packets. Additional details regarding auto-negotiation are disclosed in Breyer et al., "Switched and Fast Ethernet: How It Works and How to Use It", Ziff-Davis Press, Emeryville, California (1995), pp. 60-70, and Johnson, "Fast Ethernet: Dawn of a New Network", Prentice-Hall, Inc. (1996), pp. 158-175.

As described above, the PHY transceivers of the link partners perform auto-negotiation amongst each other to determine the best common operating mode. Each PHY transceiver then supplies any necessary capabilities information to the corresponding MAC via a management data input/output (MDIO) serial interface on the MII.

Newer PHY transceivers, such as the commercially-available Am79C875, "NetPHY™-4 LP Low Power Quad 10/100-TX/FX Ethernet Transceiver" from Advanced Micro Devices, Inc., Sunnyvale, California, integrate a plurality of PHY transceivers onto a single semiconductor chip. However, if a PHY transceiver is configured via autonegotiation, only the NID that was involved in the autonegotiation knows of the results and the other NIDs on the chip have no direct method of determining the results of the autonegotiation.

DISCLOSURE OF THE INVENTION

There is a need for a network interface in communication with a plurality of physical layer devices on a single chip to allow the results of autonegotiation to be overridden, and also to configure each MAC to agree with the results of auto-negotiation.

There is also a need for a network interface that is able to communicate auto-negotiation configuration results between a PHY device and its associated NID when there is no direct MDIO connection between the two devices.

There is also a need for a network interface that enables a plurality of physical layer devices and corresponding link partners to be configured for communicating network data by either a configuration source or through auto-negotiation.

There is also a need for a configuration source to configure the speed, type of link, link status, and pause ability of physical layer devices on a single chip.

These and other objects are attained by the present invention, where either autonegotiation or a configuration source configures physical layer devices and corresponding link partners thereby allowing network data to be transmitted and received.

According to one aspect of the present invention, a network interface is provided comprising a plurality of physical layer devices (PHYs), each configured for communicating network data to a link partner according to one of an autonegotiation protocol with the corresponding link partner and prescribed configuration information from a shared management data bus. The network interface includes a plurality of network interface devices (NIDs), each NID respectively corresponding to one of the PHYs, and each NID having a media access controller (MAC), with at least one of the NIDs configured as a master NID for communicating configuration information between the PHYs and a configuration source. A shared management data bus couples together each of the PHYs and the master NID, the shared management data bus carrying the prescribed configuration information and autonegotiation results. A configuration source is coupled to the NIDs, the configuration source being configured to receive the autonegotiation results from each PHY through the master NID and send MAC configuration information to the MAC in each of the NIDs to configure the MAC of each NID in accordance with the autonegotiation results for the corresponding physical layer device.

Another aspect of the present invention provides a method for configuring network interface devices, comprising the steps of configuring a plurality of physical layer devices (PHYs), each PHY communicating network data to a link partner according to one of an autonegotiation protocol with the corresponding link partner and prescribed configuration information from a shared management data bus. The autonegotiation results from the PHYs based on autonegotiation of the PHYs with their corresponding link partners are received at a configuration source. The individual media access controller (MAC) of each of a plurality of network interface devices (NIDs) is configured by a configuration source based on the autonegotiation results received by the configuration source from the PHYs, each NID corresponding to a respective PHY.

Additional advantages and novel features of the invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The advantages of the invention may be realized and attained by means of the instrumentalities and combinations particularly pointed out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Reference is made to the attached drawings, wherein elements having the same reference numeral designations represent like elements throughout and wherein:

Figure 1 is a block diagram of an exemplary network interface in an Ethernet (ANSI/IEEE 802.3) network according to the preferred embodiment of the present invention.

Figure 2 is a flow diagram illustrating the method for configuring of the physical layer devices and network interface devices in the network interface of Figure 1.

Figure 3 is a block diagram illustrating the format of an MDIO frame communicated between an MDIO logic block of an NID and the MDIO logic block of a PHY.

Figure 4 is a flow diagram illustrating a method for configuring MACs of non-master NIDs in accordance with embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Figure 1 is a block diagram of an exemplary network interface 10 in an Ethernet (ANSI/IEEE 802.3) network where the network interface 10 includes a plurality of physical layer devices (PHYs) 14a, 14b, 14c, and 14d. Each of the PHYs 14a, 14b, 14c, and 14d is configured for communicating network data to a link partner (not shown) according to one of an autonegotiation protocol or prescribed configuration information 30 from a shared management data bus 16. A plurality of network interface devices (NIDs) 18a, 18b, 18c, and 18d, each having a data bus interface 20a, 20b, 20c, and 20d, respectively, receive network interface configuration information 22 from a configuration source 24 via a data bus 26. The network interface configuration information 22 comprises media access controller (MAC) configuration information 28 and the prescribed configuration information 30. In addition to the data bus interface 20a, 20b, 20c, and 20d, each NID 18a, 18b, 18c, and 18d includes a media access controller (MAC) 32a, 32b, 32c, and 32d. The MAC configuration information 28 is used to configure the MACs 32a, 32b, 32c, and 32d to transmit and receive the network data across a media independent interface (MII) 34a, 34b, 34c, and 34d, respectively, to a corresponding PHY 14a, 14b, 14c, and 14d, respectively. One of the NIDs 18a, 18b, 18c, and 18d is configured as a master device (18a) for selectively supplying the prescribed configuration information 30, via the shared management data bus 16, to at least one of the PHYs 14a, 14b, 14c, and 14d.

In the preferred embodiment, the configuration source 24 is a central processing unit (CPU) and the data bus 26 is a PCI local bus. In an alternate embodiment, the configuration source 24 is an electronically erasable programmable read only memory (EEPROM) and the data bus 26 is a serial interface. In the alternate embodiment, the EEPROM downloads default network interface configuration information into the network interface devices 18a, 18b, 18c, and 18d and physical layer devices 14a, 14b, 14c, and 14d, respectively.

In an enhanced embodiment, the CPU is able to determine the operating characteristics of the NIDs 18a, 18b, 18c, and 18d and PHYs 14a, 14b, 14c, and 14d and configures each device to improve the performance for the network interface 10. Therefore if the network data error rate changes, the

CPU is able to reconfigure the network interface components to improve the performance of the network interface. For example, if too many errors are detected when a network segment is running at 100 Mb/s, the CPU can reconfigure the MAC and corresponding PHY to run at 10 Mb/s.

5 The data bus interface 20a, 20b, 20c, and 20d, for each NID 18a, 18b, 18c, and 18d, respectively, receives the network interface configuration information 28. Each of the data bus interfaces 20a, 20b, 20c, and 20d, is connected to the MACs 32a, 32b, 32c, and 32d, respectively. Each of the MACs 32a, 32b, 32c, and 32d contains at least one register 36a, 36b, 36c, and 36d, respectively, which is configured to store the MAC configuration data 28. Each media access controller 32a, 32b, 32c, and 32d is configured to transmit and receive the network data according to the MAC
10 configuration data 28.

Each of the NIDs 18a, 18b, 18c, and 18d includes management data input/output (MDIO) logic 40a, 40b, 40c, and 40d, respectively. The prescribed configuration information 28 for configuring the PHY 14a, 14b, 14c, and 14d, is received by all of the NIDs 18a, 18b, 18c, and 18d, respectively. The MDIO logic 40a of the selected or master NID 18a writes the configuration information to the PHY's 14 a-d. Hence, the selected or master NID 18a forwards the prescribed configuration
15 information 30 to common MDIO logic 42 via a management data clock/MDIO (MDC/MDIO) signal path 44. The common MDIO logic 42 receives the prescribed configuration information 30 and sends the prescribed configuration information 30 to at least one PHY 14a, 14b, 14c, and 14d.

Sharing the common MDIO logic 42 allows a plurality of PHY's to be incorporated into a
20 single chip 12, e.g., the single quad-PHY device as shown in Figure 1. Once a PHY 14a, 14b, 14c, and 14d and a NID 18a, 18b, 18c, and 18d are properly configured, network data can be transmitted and received over the MIIs 34a, 34b, 34c and 34d. In addition, each PHY 14a, 14b, 14c, and 14d is then able to transmit and receive network data to and from a network medium (not shown), which in turn allows the NID 18a-d to transmit and receive network data to and from corresponding link partners
25 (not shown).

Figure 2 illustrates the different steps used to configure the PHYs and NIDs in the network interface. The method for configuring NIDs starts with the data bus interface receiving network interface configuration information from a configuration source via a databus (step 50), where the network interface configuration information comprises the MAC configuration information and the
30 prescribed configuration information. Specifically, the MDIO logic of the NID receives the prescribed configuration information and the MAC receives the MAC configuration information. The MAC is configured according to the received MAC configuration information (step 51). The configuration source 24 sends the prescribed configuration information to a master or selected NID 18a via a data bus 26 (step 52). The prescribed configuration information is sent to common MDIO logic 42 via a
35 MDC/MDIO signal path (step 53). The MDIO logic 42 determines the address for the PHY and the address for the register for the PHY from the prescribed configuration information and then sends the

information to the appropriate PHY via a shared management databus 16 (step 54). The PHY reads the prescribed configuration information and configures the PHY accordingly (step 55).

Figure 3 illustrates an example of the format of an MDIO frame communicated between an MDIO logic block of an NID and the MDIO logic block of a PHY. Each of the serial data streams 60 starts with a preamble (32 bits) 62 comprising a known pattern of 1's and 0's. The address of the PHY (5 bits) 64 and the address of the register (5 bits) 66 for the PHY are provided. A control or read/write bit 68 is used to indicate whether the data is to be read or to be written. The turn around cycle (2 bits) 70 is the actual time that is required to reverse the direction of transmission between the PHY and the NID during half-duplex transmission. The data stream 60 also includes 16 bits of data 72.

As previously mentioned, the configuration parameters can be overridden. The configuration parameters include pause ability, link status and speed control. Pause ability is the ability of the MAC in the NID to respond to a received pause frame. Pause ability is controlled independently from the transmission of pause frames. Therefore, if the pause ability of the MAC is not enabled, the MAC will not react to a pause frame. However, if the pause ability is enabled, the MAC will finish transmitting the current frame and then stop transmitting for the amount of time indicated by the pause frame. Pause ability is enabled by either the negotiate pause ability (NPA) or the force pause ability (FPA) bits. If the FPA bit is set, pause ability is enabled regardless of the pause ability state of the link partner. If the NPA bit is set, after reset the MDIO logic 40 in the NID sends a configuration message over the shared management data bus 16 to the PHY device 14 to cause the PHY device 14 to set its internal pause ability bit. When this bit is set, during the auto-negotiation process the PHY 14 will advertise that the NID is capable of responding to pause frames. Then, if the auto-negotiation results indicate that the link partner is also capable of responding to pause frame, pause ability will be enabled in the NIDs on both ends of the link. If neither NPA nor FPA is set, pause ability is disabled and the auto-negotiation process passes this information to the link partner so that the link partner can also disable pause ability.

The present invention provides a method for configuring each MAC 32a-d to agree with the auto-negotiation results when one or more of the MACs 32a-d are unable to obtain this information directly. This issue arises when more than one MAC 32a-d is connected to a multi-phy device 12, such as that shown in Figure 1. In order to accomplish this, the MAC devices 32a-d of the present invention are flexible enough to act as either a master MAC with an MDIO interface or a non-master MAC with no MDIO connection. With a multi-PHY device 12 with a single MDIO interface 42, the CPU (configuration source in Figure 1) 24 sets up all but one MAC (32a, e.g.) so that these other MACs (32b-d, e.g.) are not MDIO masters. This is done by setting a Disable Port Manager bit in each MAC 32b-d that is not an MDIO master. The CPU 24 then uses the MDIO master MAC 32a as a path to the PHYs 14a-d. Through this path, the CPU 24 reads status information from all of the PHYs 14a-d to determine when auto-negotiation is complete. The CPU 24 reads the auto-negotiation results from all

the PHYs 14a-d, and then writes the MAC configuration information to the MAC devices 32a-d to configure the MACs 32a-d to be consistent with the PHYs 14a-d.

When the MDIO interface (40a, e.g.) of a NID (18a, e.g.) is connected to the MDIO interface 42 of a PHY device 14a, the Disable Port Manager bit of MDIO interface 40a is left in its default state, "0". When the Disable Port Manager bit is "0", the MDIO logic 40a in the NID 18a automatically retrieves the auto-negotiation results from the PHY 14a and configures the MAC 32a consistent with these results.

The Pause configuration is handled differently from the speed and duplex mode configuration, since speed and duplex mode are PHY attributes, while pause ability is a MAC feature. Although a MAC feature, the use of the Pause feature is negotiated by the PHYs 14a-d as part of the auto-negotiation process. If the CPU 24 wants to enable the PAUSE feature, the CPU 24 uses the MDIO path through the NID 18a-d to set the Pause bit in the auto-negotiation register of the corresponding PHY 14a-d and to set the "Restart Auto-negotiation" bit in the PHY's Control Register. In the case where the NID 18a-d is connected to the PHY 14a-d through the MDIO interface 42, the MDIO logic 40a-d of the NID 18a-d will automatically enable or disable Pause Ability in the NID 18a-d based on the auto-negotiation results. If the NID's MDIO interface 40a-d is not connected, the CPU 24 must retrieve the auto-negotiation result and write to the NID control registers to enable or disable Pause Ability.

To allow the NID 18a-d to act either as an MDIO master or as a device with no MDIO interface, the NIDs 18a-d have the force pause ability and negotiate pause ability control bits described earlier. When the negotiate pause ability bit is set, the MDIO logic causes the Pause bit in the external PHY 14a-d to be set as described above, then enables or disables pause ability in the NID 18a-d based on auto-negotiation results. When the force pause ability bit is set, pause ability in the NID 18a-d is enabled regardless of auto-negotiation. When both bits are cleared, the MDIO logic 40a-d in the NID 18a-d causes the Pause bit in the external PHY 14a-d to be cleared so that the auto-negotiation process will cause pause ability to be disabled in the NIDs at both ends of the network link.

Speed control can be set to any of the following: determined by the PHY, 1 Mbs, 10 Mbs, 100Mbs or 1000 Mbs.

Figure 4 depicts a basic flow chart of a method of forcing the MACs 32a-d into the appropriate configuration based on the auto-negotiation results, even when one or more of the MACs 32a-d is unable to obtain this information directly.

In step 80, the PHYs 14a-d are written to by the configuration source 24 to set up pause ability. The PHYs 14a-d are written to again to start the auto-negotiation process, in step 82. Polling of the PHYs 14a-d takes place until the status information in the PHYs 14a-d indicates that auto-negotiation is finished. The polling takes place in step 84. Upon completion of the auto-negotiation process for a PHY 14a-d, the auto-negotiation results are read from the PHY 14a-d by the configuration source 24,

88.

5